

# The Evolution of Wikipedia's Norm Network

Bradi Heaberlin<sup>1,2</sup> and Simon DeDeo<sup>1,3,4,5\*</sup>

<sup>1</sup> Program in Cognitive Science, Indiana University, 1900 E 10th St, Bloomington, IN 47406, USA; bmheaber@umail.iu.edu

<sup>2</sup> Department of Political Science, Indiana University, 1100 E 7th St, Bloomington, IN 47405, USA

<sup>3</sup> Center for Complex Networks and Systems Research, Department of Informatics, Indiana University, 919 E 10th St, Bloomington, IN 47408, USA

<sup>4</sup> Ostrom Workshop in Political Theory and Policy Analysis, 513 N Park Avenue, Bloomington, IN 47408, USA

<sup>5</sup> Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

\* Correspondence: sdedeo@indiana.edu; Tel.: +1-812-856-2855; Fax: +1-812-856-3825

Academic Editor: Emilio Ferrara

Received: 4 December 2015; Accepted: 6 April 2016; Published: 20 April 2016

**Abstract:** Social norms have traditionally been difficult to quantify. In any particular society, their sheer number and complex interdependencies often limit a system-level analysis. One exception is that of the network of norms that sustain the online Wikipedia community. We study the fifteen-year evolution of this network using the interconnected set of pages that establish, describe, and interpret the community's norms. Despite Wikipedia's reputation for *ad hoc* governance, we find that its normative evolution is highly conservative. The earliest users create norms that both dominate the network and persist over time. These core norms govern both content and interpersonal interactions using abstract principles such as neutrality, verifiability, and assume good faith. As the network grows, norm neighborhoods decouple topologically from each other, while increasing in semantic coherence. Taken together, these results suggest that the evolution of Wikipedia's norm network is akin to bureaucratic systems that predate the information age.

**Keywords:** social norms; norm networks; Wikipedia; oligarchy; bureaucracy; governance; knowledge commons

## 1. Introduction

A society's shared ideas about how one "ought" to behave govern essential features of economic and political life [1–6]. Outside of idealized game-theoretic environments, for example, economic incentives are supplemented with norms about honesty and a higher wage is possible when workers believe they ought not to cheat their employer [7]. And, while the rational structure of rules and laws is an important part of coordinating actions and desires [8], people determine the legitimacy of these solutions based on beliefs about fairness and authority. A police force without legitimacy cannot enforce the law [9,10].

Norms are also under continuous development. The modern norm against physical violence, for example, has unexpected roots and continues to evolve [11–13]. Yet, we understand far less about the history and development of norms than we do about economics or the law [14]. We often lack the data that would allow us to track the coevolution of complex, interrelated and interpretive ideas, such as honesty, fairness, and authority, the way we can track prices and monetary flows or the creation and enforcement of statutes.

Online systems, such as Wikipedia, provide new opportunities to study the development of norms over time. Along with information and code repositories at the center of the modern global economy, such as GNU/Linux, Wikipedia is a canonical example of a knowledge commons [15–18]. Knowledge commons rely on norms, rather than markets or laws, for the majority

of their governance [19,20]. On Wikipedia, editors collaborate to write encyclopedic articles in a community-managed open source environment [21,22], and they rely on social norms to standardize and govern their editing decisions [23]. Wikipedia's minute-by-minute server logs cover more than fifteen years of norm creation and evolution for a population of editors that has numbered in the tens of thousands. Norms matter on Wikipedia in ways that make it impossible for participants to ignore: it is the system of norms, rather than just laws, that dictates what content is or is not included, who participates, and what they do.

Paralleling findings in the study of rule evolution in large academic institutions [24], we expect Wikipedia's norms to play a role in the preservation of institutional memory, to be a source of both institutional stability and change, and to bear a complex relationship to the circumstances that led to their creation. Norm pages play key roles in coordinating behavior among the encyclopedia's editors [25]. Editors commonly cite norms on article talk pages in an attempt to coordinate [26], build consensus, and resolve disputes [23,27].

This study focuses on a subspace of the encyclopedia devoted to information and discussion about the norms of the encyclopedia itself. The communities associated with each of Wikipedia's 291 languages and editions have a great deal of independence to define and change the norms they use; thus, each can follow a different evolutionary trajectory. Here, we focus solely on norms in the English-language Wikipedia. We study the evolution of these norms using a subset of tightly-linked pages that establish, describe, and interpret them. These pages, along with the relationships between them, allow us to quantify how editors describe expectations for behavior and, consequently, how they create and reinterpret the norms of their community.

We focus on the links between norm pages. Online link formation occurs for a variety of reasons [28], including strategic association by the individual making the citation [29]. In the case of Wikipedia, links between pages in the encyclopedia "main space" encode information about semantic relationships [30,31] and the relative importance of pages [32,33]. Extending these analyses to the norm pages of the encyclopedia allows us to see how norms are described, justified, and explained by reference to other norms. Our use of this network parallels studies of citations in legal systems; researchers use legal citations to track influence via precedence [34] and legitimation [35], as well as the prestige of the cited [35,36]. The parallel to legal citations is not exact: the pages within Wikipedia's norm network are not (usually) created in response to a particular event, as in a court case, but rather in response to a perceived need; pages can be created by any user, rather than a particular judge or court; and pages can be retrospectively edited (leading, for example, to the potential for graph cycles when new links are introduced).

This network perspective allows us to go beyond the tracking of a single behavior over time (a common approach in studies of cultural evolution [37]) to look at the evolution of relationships between hundreds, and even thousands, of distinct ideas. We use these data to ask three critical questions. In a system where norms are constantly being discussed and created, how and when do some norms come to dominate over others? What types of behavior do they govern? Additionally, how do those core norms evolve over time?

The answers are surprising. While some accounts of Wikipedia stress its flexibility and the *ad hoc* nature of its governance [38–40], we find that Wikipedia's normative evolution is highly conservative. Norms that dominate the system in Wikipedia's later years were created early, when the population was much smaller. These core norms tell editors how to write and format articles; they also describe how to collaborate with others when faced with disagreements and even heated arguments. To do this, the core norms reference universal, rationalized principles, such as neutrality, verifiability, civility, and consensus. Over time, the network neighborhoods of these norms decouple topologically. As they do so, their internal semantic coherence rises, as measured using a topic model of the page text. Wikipedia's abstract core norms and decoupling process show that it adopts an "institutionalized organization" structure akin to bureaucratic systems that predate the information age [41].

## 2. Methods

To gather data on the network of norms on Wikipedia, we spider links within the “namespace” reserved for (among other things) policies, guidelines, processes, and discussion. These pages can be identified because they carry the special prefix “Wikipedia:” or “WP:”. Network nodes are pages. Directed edges between pages occur when one page links to another via at least one hyperlink that meets our filtering criteria; these links are found by parsing the raw HTML of each page and excluding standard navigational templates and lists. Our network is thus both directed and unweighted. We begin our spidering at the (arbitrarily selected) norm page “Assume good faith”. Details of the spidering process, hyperlink filters and our post-processing of links between pages appear in Appendix A; both the raw data and our processed network are freely available online [42].

Editors classify pages in the namespace by adding tags; these tags include, most notably, “policy”, “guideline”, and “essay”, among others. When we download page text, we also record these categorizations. These categorizations describe gradated levels of expectations for adherence [43]. In automatically-included “template” text, policies are described as “widely accepted standards” that “all editors should normally follow” [44], guidelines as “generally accepted standards” that “editors should attempt to follow” and for which “occasional exceptions may apply” [45], while essays provide “advice or opinions”: “[s]ome essays represent widespread norms,” while “others only represent minority viewpoints” [46]. A fourth category is the “proposal”, which describes potential policies and guidelines “still ... in development, under discussion, or in the process of gathering consensus for adoption” [47].

Previous analysis of Wikipedia’s policy environment has emphasized the many, often overlapping, functions that norms play in the encyclopedia, such as policies that both attempt to control un-permitted use of copyrighted material and to establish legitimacy through the use of legal diction and grammar [25]. In the current study, we consider a complementary classification system that focuses on the types of interactions the norms govern, rather than their functions. We propose three distinct norm categories based on, and extending, pre-existing classification of the norms that govern natural [19] and knowledge commons [20].

Norms may attempt to regulate content creation (“user-content” norms) and interactions between users (“user-user” norms). In addition, norms may attempt to define a more formal administrative structure with distinct roles, duties, and expectations for admins (“user-admin” norms). The two authors of this paper independently categorized a random sample of forty pages using this scheme, and we calculated inter-coder reliability using Cohen’s kappa [48].

For our semantic analysis, we include all text, except that found in special boxes whose text is replicated by template across multiple pages. To build our distribution over one-grams, we normalize all text to lowercase, merge hyphenated words (“error-correction” to “errorcorrection”), and drop punctuation (“don’t” to “dont”). We do neither stemming nor spelling correction.

A critical external variable is the number of active users on the encyclopedia at any point in time. Following [49], we define an active user as one who has made five or more edits within a month; these statistics are publicly maintained at [50].

### 2.1. Centrality and Attention Measures

The pages in our corpus are created to explain the norms of Wikipedia to editors and influence their interactions with the encyclopedia’s editing community and content. Users navigate the system of norms as a network structure and consequently encounter some pages more than others.

We measure this using eigenvector centrality (EC), which quantifies the importance of a page based on its overall accessibility within the network. The EC of a page is the probability of happening across a page during a random walk; equivalent to the PageRank algorithm, it is used in the behavioral sciences to identify consensus on dominance and power [51]. We set  $\epsilon$ , the probability of a random jump, to 0.15.

We expect some pages to become highly central to the network, while others remain largely peripheral. We quantify the inequality of the system using the Gini coefficient (GC). GC varies between zero (perfect equality; all pages have equal EC) and one (one page has a high EC; all other pages have the same low value). GC is widely used in economics to measure income inequality. Here, it provides a global measure of the extent to which a system is dominated by a few norms. As a dimensionless quantity, it allows researchers to compare this system to others that might be the subject of later research.

Because we are interested in the ways in which the norm citation network evolves and the role that norms play in the context of this structure, EC is an ideal measure of a norm's importance. In addition to quantifying structural importance, however, we expect EC to correlate with, and to predict, behavioral measures of the attention a page receives. To measure the relationship between centrality and behavioral measures of attention, we track page view data (from Wikipedia's server logs made available by StatsGrok [52], see Appendix B), the total number of edits a page has received, the number of edits on its associated talk page, and the number of editors who have edited the page. We perform a multivariate linear regression on these attention measures, along with page age and page size (in bytes) as predictors of a page's EC (see Appendix C).

## 2.2. Influence and Overlap

An important feature of the norm network is the sphere of influence: the pages that rely on any particular page for context.

Consider, for example, the norm page "Neutral Point of View" (NPOV), a page urging editors to describe article subjects without taking sides. A page that links to NPOV relates its own subject to NPOV in some fashion. For example, among many pages that link to NPOV is "Propaganda", an essay urging editors to be wary of using propaganda outlets of authoritarian governments. The Propaganda page links to the NPOV page in order to define the notion of "undue weight"; NPOV's content can thus be said to influence the interpretation of what is found on Propaganda.

Influence is distinct from centrality; centrality measures the extent to which pages link to the page in question. Conversely, influence measures the extent to which the content of that page influences other pages. In our formalism, a node  $p$  can be understood to *influence* a node  $q$  when  $q$  links to  $p$ . Influence need not be direct, however:  $p$  can influence  $q$  if  $q$  links to  $r$  and  $r$  links to  $p$ . To measure the non-local influence, we consider random walks on the direction-reversed network.

More formally, placing a random-walker at node  $p$ , we allow her to take  $n$  steps from this starting point along the direction-reversed network; we write the resulting probability distribution over the final position as  $p_i$ , the probability of the walker ending up at node  $i$ . The distribution  $p_i$  defines the influence that  $p$  has on  $i$ .

To quantify the distance between two nodes, we then consider the influence *overlap* between two arbitrary nodes  $p$  and  $q$ . Overlap quantifies the extent to which two random walkers, beginning at these nodes, will tend to visit the same pages. If  $p_i$  and  $q_i$  are the probability distributions associated with the influence of node  $p$  and  $q$ , then overlap is defined as:

$$O(p, q) = \frac{\sum_{i=1}^N p_i q_i}{\left( \left[ \sum_{i=1}^N p_i^2 \right] \left[ \sum_{i=1}^N q_i^2 \right] \right)^{1/2}} \quad (1)$$

For multiple pages, we can compute the average pairwise overlap simply by averaging the overlap between all possible pairs within the set.

High overlap between  $p$  and  $q$  indicates that two pages influence a large number of common nodes. When  $n$  goes to infinity, the random walkers converge to the stationary distribution, and the overlap is one; conversely, when  $n$  is small, random walkers have less time to encounter each other. We take  $n$  equal to five, larger than the average shortest path (roughly three, in our network), so that nodes are potentially reachable, but much less than the convergence time to the stationary distribution.

Overlap can be thought of as a measure of the separation of spheres of influence. It invokes only local mechanisms: users traveling from one page to another by the links that connect them. This is in contrast to a measure, such as shortest paths, which is computationally expensive and requires detailed, global knowledge of the network link-structure. In general, for example, the number of nodes an algorithm needs to visit in order to determine the shortest path between two nodes will usually be much larger than the length of the final path.

Both influence and overlap require us to specify particular nodes of interest; we focus in this work on pairs of high-EC pages, or core norms.

### 2.3. Semantic Coherence

We consider the semantic relationships between pages. This provides a notion of relatedness that is distinct from how norms connect via hyperlinks. To do this, we do topic-modeling (latent Dirichlet allocation [53]) on the one-grams of the visible, human-readable text on each page. Topic models allow us to represent short texts even when they draw from a rich vocabulary: topics coarse-grain the underlying distributions over words.

With the resulting topic model, we can then compute the semantic distance between all pairs of pages using the Jensen–Shannon distance (JSD), a measure that quantifies the distinguishability of two distributions [54]. This gives us a weighted semantic network that we can compare to the network of hyperlinks between pages. In particular, we can compute the *semantic coherence*: the Pearson correlation between  $p_i$  (the influence of node  $p$  on node  $i$ ) and the negative JSD from node  $p$  to node  $i$ ,  $J_i$ . When nodes that are closely related topologically are also closely related semantically (JSD low), the coherence is high.

### 2.4. Community Detection

We expect the links that editors make at the local level to give rise to distinct clusters, or *norm bundles*, at the global level. We use the Louvain community detection algorithm [55] to detect clustering among the nodes in the network. The Louvain algorithm maximizes the modularity at each local partition of the network. The algorithm first assigns each node  $i$  to a different cluster, then computes the potential modularity gain to  $i$  for joining the cluster of its neighbor node  $j$ . Each  $i$  will join the cluster of  $j$  when the merge offers the highest positive modularity gain. If there is no possible gain in modularity,  $i$  remains in its initial cluster.

## 3. Results

At first, Wikipedia's population underwent exponential growth. In mid-2007, however, population growth stalled and entered a period of secular decline [49]; see Figure 1. Over the course of this rapid growth and longer timescale decay, users created a large number of pages establishing, describing, and interpreting community norms. Our analysis finds a total of 1976 pages associated with norms. There are 17,235 edges between these nodes; the network density, 0.0044, is of the same order of magnitude as those seen for academic citation networks [56]; 1872 (95%) of these pages are linked together in a giant component.


There are a total of 56 pages classified as policy and 113 marked as guideline; for concision, we refer to pages of both types as “policy”. The majority of non-policy pages (1807) are classified as “essays” (1255), followed by “proposals” (182) (suggestions either rejected by the community or under discussion), and “humor” pages similar to essays, but taking a more irreverent tone (125).

We were able to achieve good, but not perfect, agreement in categorizing pages as user-content, user-user, or user-admin norms. Our categorization agreement rate was 75% over forty randomly-selected pages. This is well above chance ( $p \ll 10^{-3}$ ), with Cohen's  $\kappa$  value, of 0.59 indicating “moderate” agreement [57]. We disagreed, for example, on “Editors\_should\_be\_logged-in\_users\_(failed\_proposal)” (user-user *vs.* user-content) and “Paid\_editor's\_bill\_of\_rights” (user-user *vs.* user-admin). In the same



sample of forty random pages, we encountered only one that we believed was not a norm, giving an approximate precision rate of 97.5%.

### 3.1. Network Construction




**Figure 1.** Cumulative growth in policy (red/solid line) and non-policy (green/dashed line) pages, overlaid on active population (blue/dotted line). Policy creation precedes the arrival of the majority of users, while the creation of non-policy pages, usually in the form of essay and commentary, lags the growth in population.

Most policy pages appear before the bulk of the population arrives: over half the policy pages were created by May 2005, before the population reached 20% of its maximum. By the time the population did reach its maximum, in March of 2007, over 80% of the policy pages had already been created. By contrast, the creation of non-policy pages in the form of essays and commentary lags population growth. When the population reached its March 2007 maximum, less than one-third of the non-policy pages were in place. It was not until a year later that half of the policy pages were in place. This is shown in Figure 1.

Eigenvector centrality leads to a distinct hierarchy of pages, with some gaining a significant fraction of the overall centrality in the system. This is shown in Appendix D, Figure D.1, broken out by four main page categories—policies, guidelines, essays, and proposals. Policies and guidelines dominate the system by centrality. Our centrality measure correlates with all of the behavioral measures of attention we consider (see Appendix B, Table B.1).

The hierarchy is established early and yet is remarkably stable over the lifetime of the system. The Pearson correlation between the eigenvector centrality of nodes in 2001 and their final values in 2015 is 0.87; year to year, it is always greater than 0.9. The growth in nodes' in-degree is roughly multiplicative; for nodes with degree less than one-hundred (93% of the total network), the growth rate is, on average,  $12.7\% \pm 0.3\%$  from one year to the next. There is some evidence for super-multiplicative returns to scale; the yearly growth rate for pages with in-degree less than ten is only  $10.6\% \pm 0.4\%$ .

All of this means that, as new pages enter the system, they fail to gain the prominence of the early core norms. This leads to an increase in overall network inequality, shown in Figure 2.



**Figure 2.** Evolution of the Gini coefficient over time. As new pages enter the system, overall network inequality increases, stabilizing in 2008.

In short, policy growth precedes population growth. Policies have far greater centrality in the network than other page types. Centrality in the network is unequally distributed and becomes less equal over time.

### 3.2. Core Norms

Table 1 lists the top twenty pages in our network. These core norms govern a range of behaviors, including user-content actions (write articles from a neutral point of view, #1; include only verifiable information, #2; and reliable sources, #3), user-user actions (find consensus, #6; assume good faith, #11; be civil, #16; do not “edit war”, #19), and user-admin relationships involving specially-defined roles (blocking policy, #13; the arbitration committee, #17). In some cases, a norm spans multiple classes; “What Wikipedia is not”, for example, includes both “Wikipedia is not a dictionary” (a norm on the nature of the content to be included) and “Wikipedia is not a battleground” (a norm on how users should interact with each other).

**Table 1.** Core norms. Top twenty pages, by eigenvector centrality, in 2015. All are either policy or guideline pages, and all were in place by the end of 2006. The majority of these core norms were created before 2004, when the population was less than 3% of its peak.

Rank	Name	Classification	Creation Date
1	Neutral_point_of_view	User-content	24 December 2001
2	Verifiability	User-content	2 August 2003
3	Identifying_reliable_sources	User-content	28 February 2005
4	What_Wikipedia_is_not	User-user/user-content	24 September 2001
5	Biographies_of_living_persons	User-content	17 December 2005
6	Consensus	User-user	11 July 2004
7	Policies_and_guidelines	User-user/user-content	1 November 2001
8	Administrators	User-admin	16 May 2001
9	No_original_research	User-content	21 December 2003
10	Citing_sources	User-content	19 April 2002
11	Assume_good_faith	User-user	3 March 2004
12	Notability	User-content	7 September 2006
13	Blocking_policy	User-admin	8 June 2003
14	Dispute_resolution	User-user/user-admin	12 January 2004
15	Redirect	User-content	25 February 2002
16	Civility	User-user	5 February 2004
17	Arbitration_Committee	User-admin	16 January 2004
18	Vandalism	User-content	29 March 2002
19	Edit_warring	User-user	26 April 2003
20	Talk_page_guidelines	User-user	15 April 2005

All of these core norms were created early in the system's history. The majority were created before 2004, when the population was less than 3% of the March 2007 peak. The earliest members of the community first defined and articulated its core norms.

It is important to note that while the most important norms are those that are created early, not all of the pages created early become, or remain, central to the network. This is shown visually in Appendix C, Figure C.1; there are many old pages that never grew to importance and that have ECs comparable to the youngest pages. Because of this, page age alone is not a significant predictor of eigenvector centrality. We confirm this with a multivariate linear regression (see Table C.1). The number of editors is a strong predictor; not only do high EC pages attract a large number of unique editors, but there are few low-EC pages that do.

### 3.3. Overlap and Semantic Coherence

Over the course of network construction, core norms are drawn apart topologically. At the same time, the semantic coherence of their neighborhoods rises.


Figure 3 shows the average pairwise overlap between the top ten pages in our network (since some norms are created later, the number of norms in this final set is lower early on). Early in the system history, when the network is small, overlap is very high. The creation of new pages leads to a rapid decline in overlap; even in 2006, when all core norms are in place, the overlap continues to decline. Figure 3 also shows the evolution of semantic coherence, which rises rapidly and stabilizes early.

Network growth could have been imagined to drive a knitting together of distinct principles. Instead, the opposite happens: core norms slowly draw apart as page creation leads to distinct spheres of influence. Rather than nucleating around a set of densely-connected core principles, the norm network continues to condense around multiple points.

We note that the local clustering coefficient, a measure of the extent to which two nodes, linked to the same node, tend to also link together, remains essentially constant over the span of the data (see Appendix E, Figure E.1). The ways in which editors link together small groups of pages changes



little while their cumulative effect produces large and lasting changes both in attention inequality and page overlap.




**Figure 3.** Evolution of influence overlap among the core norms (top twenty norms by eigenvector centrality) over time (solid line, labeled). In terms of the pages they influence, core norms draw apart over time, stabilizing in 2008. At the same time, semantic coherence (dashed line, labeled) increases. Neighborhoods become topologically distinct, but internally coherent.

### 3.4. Emergent Clusters

The connected component of network, containing 95% of all nodes, partitions into 10 clusters. In Table 2, we describe the top nine, which together nearly all of the giant component. By inspecting the top ten nodes in each cluster, we classify them into user-content, user-user, and user-admin norms (see Table F.2). A force-directed layout (ForceAtlas2, implemented in Gephi [58]) allows us to visualize the norm network and the topological relationships between its emergent groups (see Figure 4).

**Table 2.** Top nine Louvain clusters, by number of nodes. Communities fall into three classifications (user-user, user-content, user-administration), based on the interactions they govern; we determine these labels by inspecting the top ten nodes by centrality within each cluster.

Rank	Fraction of System	Classification	Topic
1	24.8%	User-Content	Article Quality
2	22.9%	User-User	Collaboration
3	17.1%	User-Administration	Administrators
4	14.7%	User-Content	Formatting Articles
5	10.5%	User-Content	Content Policies
6	5.4%	User-User	Wiki-larping
7	2.0%	User-Content	Page Templates
8	1.3%	User-User/User-Content	Experts and Credentials
9	1.0%	User-User	Humor



**Figure 4.** The topology of the norm network is organized around five central clusters, found using the Louvain algorithm. Cluster themes are based on a sample of high-eigenvector centrality (EC) nodes in each cluster and confirmed by reference to a topic model of word usage. Left panel: full network, with cluster membership indicated by color. Right panel: cluster structure. Each node is a Louvain cluster, and node size indicates cluster size by number of pages. Edge weights are defined as the fraction of the origin cluster's out-links that link to each other cluster (self-loops are not shown).

The five largest clusters comprise roughly 90% of the network. The Article Quality cluster includes nodes such as Neutral Point of View, Verifiability, and Reliable Sources, governing how articles should be written. The Collaboration cluster includes pages on Consensus, Assume Good Faith, and Edit Warring, describing policies and norms associated with interpersonal interaction. The Administrators cluster contains pages relevant to administrative actions, such as the Blocking Policy and the Arbitration Committee. The Formatting cluster contains articles such as Manual of Style, Article Titles, and Disambiguation. Additionally, the Content Policies cluster contains articles on copyrights, copyright violations, and policies on image use and use of non-free content. The remaining clusters include a small group of articles on page templates; one on the role of experts of Wikipedia; two groups of humor pages (Wiki-larping, a humorous take on Wikipedia as if it were a Dungeons and Dragons game, and a cluster of pages, including “Bad Jokes and Other Deleted Nonsense”).

Each of the top nine clusters is associated with a distinct topic in our topic model (see Appendix F, Table F.1); while the article quality cluster is the largest by node number, the topic associated with the collaboration cluster dominates the system by word. Even task-based norms appear to draw on the semantics of interpersonal cooperation.

#### 4. Discussion

The most influential pages in the norm network are also the earliest to be created. A Matthew effect [59] appears to operate for social norms, where later additions to the network do not grow in influence quickly enough to destabilize the hierarchy. Why are there no normative revolutions on Wikipedia?

Perhaps the earliest users know best: their policies work well and are simply adopted by those who come later; or, later users may join precisely because they subscribe to the norms that have already

been articulated. Users who disagree with these norms may find that reinterpretation, rather than replacement, is a more effective response given the disproportionate allocation of attention to early pages.

The fact that core norms are created so early means that a relatively small number of users set them in place. This group may have created norms that meet their own needs, but not the needs of those who arrive later. For example, if early users are predominantly university students with flexible working hours, for example, they may develop norms that implicitly rely on the possibility of responding to criticism in short, rapid bursts. If later arrivals do not have the same flexibility, but the norms persist, they will find themselves at a relative disadvantage in conflicts that arise, even if the amount of effort they devote to the system each week is the same.

Recent work [60] has suggested that early users later form an oligarchy that monopolizes power, subverts democratic control, and comes into increasing conflict with the larger collective. If this is true, the enduring centrality of their own interests in the norm network may be a source of power.

Alternatively, the influence of a small group of editors may persist via the core norms despite a gradual decentralization of power within the encyclopedia. One ethnographic account of Wikipedia's editing community [61] suggests that a group of "old-timers" brings important social norms with them into the encyclopedia's increasingly local governance structures, such as WikiProject communities. Our findings show that the structure of the norm network is, by measures of page count, clustering, core norm overlap, and semantic coherence, largely stable by 2008. Thus, the core norms and global norm structure analyzed here may provide an early foundation of norms for small, decentralized communities that form later in the encyclopedia's development.

Much of Wikipedia's network simply coordinates technical practices, such as article naming conventions. The most important norms, however, attempt to rationalize the system around universal concepts, such as neutrality, verifiability, consensus, and civility. An important insight comes from a theory of bureaucracy and institutionalized organization developed by Meyer and Rowan (1977 [41]). They propose that norms such as these can function as institutional myths that make the system appear legitimate and less *ad hoc*, by connecting it to a rational framework.

Page creation continues to grow long after the core norms are already in place. What happens when editors continue to develop and refine this network?

Meyer and Rowan's theory predicts the phenomenon of decoupling, driven by the emergence of inconsistencies between different myths. The essay *Civil\_POV\_pushing*, for example, describes how some users may be able to violate the neutrality norm by strict adherence to norms of civility. In Meyer and Rowan's theory, pages like these, that attempt to resolve inconsistencies between myths, will be rare. Myths will instead tend to decouple from each other over time.

Our quantitative findings are consistent with this prediction. As the system grows, the creation of norm-spanning pages, such as *Civil\_POV\_pushing*, are rare and insufficient to prevent the neighborhoods of the core norms drawing apart into separate spheres of influence with high internal semantic coherence. In successful systems, decoupling is also expected to happen not only between myths, but between these myths and actual practice, a phenomenon pointed to by the existence of the page "*Ignore\_all\_rules*" ("if a rule prevents you from improving Wikipedia, ignore it").

Our findings are also consistent with Meyer and Rowan's second major prediction: that systems become increasingly reliant on a logic of good faith rather than following procedure. Not only is "Assume good faith" itself a core norm, but the associated topic dominates the system as a whole.

The norm network we study here is the culmination of over thirty thousand edits. We analyze the development of this system over time via the editing community's collective decisions and their allocation of attention within the network. While this method tells us a great deal about the collective process of norm creation, we do not know how individual editors understand the relationships between norms or use them to guide how they edit and interact with others. Rather than memorize the complex network in its entirety, an editor may coarse-grain its properties to form his or her own mental representation of the encyclopedia's normative structure. Editors' mental representations might then

inform their linking and editing behaviors, creating a feedback loop between the representation and the norm network as a whole.

## 5. Conclusions

Norms are a crucial unit of cultural evolution, and they gain meaning and force from the relationships that connect them. Our work here has studied the evolution, over fifteen years, of the interdependent network of norms at the center of Wikipedia.

The evolution of this network is a remarkably conservative process. Early features are maintained, and in some cases even amplified, over the course of the network's development. Our findings are consistent with the "iron law" of oligarchy in peer-production systems; they also complement accounts of gradual decentralization in Wikipedia's governance structure.

The encyclopedia's core norms address universal principles, such as neutrality, verifiability, civility, and consensus. The ambiguity and interpretability of these abstract concepts may drive them to decouple from each other over time. Wikipedia is a paradigmatic example of a 21st Century knowledge commons. Yet, its core norms play a structural role analogous to the institutional myths of rationalized 20th Century bureaucracies.

**Acknowledgments:** We thank John Miller (Carnegie Mellon), Stephen Benard (Indiana University) and Cris Moore (Santa Fe Institute) for helpful discussions, and the Santa Fe Institute for their hospitality when this work was begun. Bradi Heaberlin was supported by the Research Experience for Undergraduates program at the Santa Fe Institute under National Science Foundation Award #ACI-1358567, by the Cox Research Scholarship Program and by the Indiana University Science, Technology and Research Scholars (STARS) program. Data used in this analysis are available online [42].

**Author Contributions:** The authors jointly designed the research concept, gathered the data, and conducted the analyses. Both authors jointly wrote the paper.

**Conflicts of Interest:** The authors declare no conflict of interest.

## Appendix A. Corpus Construction

As described in the main text, we build our corpus by spidering outward from the page “Assume good faith”, following all links in the Wikipedia namespace to build a directed, unweighted network. Not all pages within the namespace are normative, however. After completing the spidering process, we remove pages that are solely lists (e.g., the pages “List of guidelines” or “Lists of protected pages”) that describe “projects”, or other initiatives focused on outreach (e.g., “Wikipedia Loves Libraries”), or on adding a certain kind of content to the encyclopedia (e.g., “WikiProject Libertarianism”), or that serve as noticeboards (e.g., the “Village pump”, “Media copyright questions”), with filters on both page titles and editor-assigned categories.

Many page names have synonyms (e.g., “AGF” redirects to “Assume good faith”); we merge synonyms. Not all links between pages indicate a deliberate decision to connect one norm to another. Many pages, for example, contain “boxes”, small code snippets that categorize pages or provide navigation indices to similar norms. These boxes can be created by a single command and are replicated across multiple pages; we do not include out-bound links found in these boxes. We do not count multiple links between pages; our edges are unweighted; a directed edge between A and B refers only to the presence of at least one link from A to B. Pages sometimes have internal links; we drop all self-edges. Our spidering includes only pages that existed on 12:00:00 UTC, 20 August 2015.


## Appendix B. Relationship between Eigenvector Centrality and Attention Measures

To compare the norm network structure and user attention, we measure the correlation between the centrality of a norm page and the percent of the network’s page views that the norm accumulates over a 31-day period (July 2015). We find a moderate correlation [62],  $r = 0.32$ , between EC and page views. (The distribution of EC and page view values is slightly non-linear. We conduct a power-law fit and find that  $\alpha = 1.42 \pm 0.02$ . Consequently, a page that doubles its EC more than doubles its share of the network’s page views. For simplicity in this analysis, we present the linear correlations.)

EC correlates significantly with all behavioral attention measures we consider; not just page views, but number of edits, number of talk page edits and number of editors; see Table B.1.

**Table B.1.** Correlation of eigenvector centrality with behavioral measures of attention.

Attention Measure	$r$	$p$ Value
Page views	0.32	$<10^{-3}$
Number of edits	0.70	$<10^{-3}$
Number of talk page edits	0.63	$<10^{-3}$
Number of editors	0.72	$<10^{-3}$



**Figure B.1.** The relationship between EC of a page and the percent of the network’s page views it accumulates.

### Appendix C. Regression on Age and Edits

To see how a page’s intrinsic properties affect its eigenvector centrality (EC) in the final network, we performed a multivariate regression with page age, number of page edits, number of talk page edits, number of editors and page size (in bytes) as predictors of EC. Including pageviews as a predictor does not significantly improve  $R^2$ ; we leave it out of our regression model.

We normalized all predictors by z-score to allow comparison between coefficients. We considered two relationships between EC and predictor variables: a linear model and a logistic model. We found the linear model has lower mean-squared error and report the coefficients in Table C.1.


**Table C.1.** Coefficients of a multivariate linear regression for page eigenvector centrality (EC). The  $R^2$  for the fit, including all predictors, is 0.57.

Predictor	Coefficient $\times 10^{-5}$	$p$ Value
Number of editors	$95 \pm 6$	$<10^{-3}$
Number of talk edits	$46 \pm 3$	$<10^{-3}$
Page size	$2 \pm 2$	n.d.
Age	$2 \pm 2$	n.d.
Number of edits	$-30 \pm 7$	$<10^{-3}$

As noted in the main text, our results show that age is a weak predictor of EC, once other variables are included. The number of unique editors is a very strong predictor, as is the number of edits to the talk page. Figure C.1 shows the distinct effect of page age and number of editors. While the most important pages are also the oldest, there are many old pages that are not important at all; the skewed distribution of eigenvector centrality means that this signal is largely washed out in a simple linear model that does not take into account the increasing variance. To reach the top 1% in EC, you must be old; but to be old is not enough.


By contrast, pages with many editors tend to be high-EC, and there are very few pages with many editors that are not also high in EC. High-EC pages not only attract more page views (see Section B), but also more editors. Interestingly, the total number of edits has a negative coefficient in our regression; while there is a strong positive correlation between the number of editors and the number of edits, there are a number of low-EC pages with many edits by a small number of people (e.g., an essay, written in many stages, by a single author, that never gains traction).






**Figure C.1.** Important pages are old, but not all old pages are important. Left panel: page age (from the end of our data, in August 2015) *vs.* eigenvector centrality; “core norms” (top twenty pages by EC) are marked by a lower bound in EC and a lower bound in age. While the very top pages in the hierarchy are all old (in the top-right region), there are many old pages that have eigenvector centrality comparable to much younger pages. Right panel: number of (unique) editors on the page *vs.* eigenvector centrality. A much tighter correlation shows that pages that attract many unique editors have higher EC. When both effects are taken into account in a simple linear regression model, the number of editors dominates.

#### Appendix D. Combined Scree Plot



**Figure D.1.** Ranked eigenvector centrality for pages, broken out by page category. Policy (blue diamond) and guideline (red plus) pages dominate the system. More interpretive essays (green squares; includes humor and related pages), the most common by number, appear at lower relative rank; the highest ranked essay, for example, has lower centrality than the 10th ranked policy. Proposals, failed or current (grey triangles), are the lowest ranked of all.



**Figure D.2.** Eigenvector centrality for all the pages in our data, ordered by rank. Major divisions (see text) are marked by vertical lines.


Figure D.1 shows the rank distribution of EC, page by page, broken out by page class. Defining  $E_i$  as the eigenvector centrality of the  $i$ -th ranked norm allows us to define the break size,  $E_i - E_{i+1}$ , between this norm and the next. Ranking break-sizes allow us to note positions where the remainder of the norms in the system have significantly lower EC. Constraining breaks to be greater than five pages apart leads to the top five divisions shown in Figure D.2. In the main paper, we list nodes up to the third break-point.

## Appendix E. Local Clustering Coefficient

Our work here focuses on the evolution of global network properties, such as eigenvector centrality, overlap and semantic coherence, that cannot be known by breaking the graph into subgraphs. It is interesting to consider more local measures, however, since these are likely to be under far greater direct user control. The example we consider here is average local clustering, defined as:

$$\kappa = \sum_{i \in G} \frac{\sum_{j,k \in \mathcal{N}(i)} \delta_{jk}}{|\mathcal{N}(i)|(|\mathcal{N}(i)| - 1)} \quad (2)$$

or, in other words, the number of edges connecting nodes in the neighborhood of  $i$ , as a fraction of the total number of possible connections between those neighbors. If individuals have a tendency to connect up the network when they create a new node, by linking together nodes it links to, this will tend to increase the clustering. Figure E.1 shows this over time. Despite large changes in both population and network size, clustering remains surprisingly constant, at around one-third.



**Figure E.1.** The average local clustering coefficient, as a function of time. Despite large-scale changes in overall network properties, this local property remains remarkably constant.

## Appendix F. Clusters and Topic Modeling

For our base model with  $k = 20$  topics, Table F.1 shows the top twenty representative words for each topic; in this table we drop the word “wikipedia”, plurals (except the word “wikipedias”) and date/time terms (“january”, “utc”, *etc.*). We use Jason Adams’ software package “lda-ruby” package (<https://github.com/ealdent/lda-ruby>), a ruby wrapper for the C code of David M. Blei; this code estimates model parameters using a variational Expectation Maximization algorithm (<http://www.cs.princeton.edu/~blei/lda-c/> [53]). In Table F.1, we show the topics, and their associated words, ordered by the topic’s (word-level) prevalence within the encyclopedia.

For each page, we can compute a distribution over topics; this is just the average of the word-level distributions. By averaging these topic distributions over pages, we can compute the topic distribution for each Louvain community (Collaboration, Article Quality, *etc.*). It turns out that each of the top eight communities has a different most-common topic. This allows us to associate some of the topics we find with a particular cluster, and we list this correspondence in column three of Table F.1. Inspection of the representative words for these eight topics provides complementary evidence in favor of the community labels, which were previously chosen by manual inspection of the top ten pages by eigenvector centrality (Table F.2).

**Table F.1.** Representative one-grams from each of the topics in our  $k = 20$  topic model, ranked by the weighted fraction of words assigned. The top nine Louvain clusters are each dominated by a unique topic.

Rank	Fraction	Louvain Community	Representative Words
1	11.4%	Collaboration	editor, edit, dont, good, people, make, editing, policy, page, talk, time, article, faith, point, policies, encyclopedia, consensus, community, personal, user
2	8.67%	Article Quality	source, reliable, article, material, information, research, primary, view, original, editors, subject, published, secondary, policy, neutral, point, scientific, content, topic, claims
3	8.56%	—	article, deletion, page, deleted, discussion, content, delete, speedy, talk, tag, subject, information, policy, user, guidelines, criteria, notability, afd, time, essay
4	8.26%	Experts and Credentials	article, information, content, encyclopedia, editors, people, wikipedias, subject, featured, quality, good, list, topic, readers, time, work, project, knowledge, number, lead
5	6.65%	—	consensus, policy, discussion, community, process, committee, arbitration, editors, administrator, user, request, policies, admin, block, dispute, page, wikimedia, proposal, information, made
6	5.80%	Formatting Articles	article, names, title, page, english, disambiguation, naming, redirect, conventions, common, term, style, citation, word, language, topic, book, usage, examples, cases
7	5.69%	Administrators	user, edit, page, vandalism, account, ip, editing, talk, editors, bot, address, protection, administrators, userboxes, username, blocked, block, request, sock, template
8	5.36%	—	notable, article, notability, list, sources, coverage, criteria, information, subject, reliable, emoji, guideline, film, event, university, significant, general, topic, independent, inclusion
9	5.03%	—	page, link, text, image, file, wikimedia, search, commons, web, information, external, software, content, article, site, add, click, wiki, edit, make
10	4.38%	—	talk, edit, page, user, war, im, article, dont, people, time, contris, good, contributions, back, long, list, things, make, day, ive
11	4.04%	Content Policies	copyright, image, public, nonfree, free, work, content, license, domain, law, fair, article, copyrighted, published, states, pma, united, subject, permission, media
12	3.81%	—	page, talk, template, namespace, user, link, article, text, category, section, special, edit, title, list, signature, ut, mediawiki, redirect, move, navbox
13	3.28%	Humor	list, chart, people, united, war, town, world, man, england, states, british, top, hot, songs, women, city, ireland, music, number, death
14	3.04%	—	category, day, categories, article, tip, stub, list, page, people, categorization, main, link, year, created, red, featured, create, template, sort, subcategories
15	2.97%	Wiki-larping	people, user, time, status, wikidragon, truth, wikifauna, wikipuma, credentials, names, editathon, work, turkish, years, page, make, real, history, group, greek
16	2.83%	—	support, oppose, policy, people, user, proposal, talk, userboxes, dont, image, offensive, pov, namespace, page, content, article, censorship, vote, npov, agree
17	2.79%	—	ban, topic, editing, indefinite, talk, sanctions, article, page, user, edit, discussion, banned, paid, related, editor, contris, interest, coi, community, broadly
18	2.55%	—	quotation, style, citing, punctuation, american, mos, ads, dash, manual, inactive, en, english, issue, sentence, dashes, election, text, space, british, jumped
19	2.41%	Page Templates	text, template, page, line, article, gt, section, lt, enforcement, table, footnote, law, summary, infobox, style, agencies, synth, color, work, data
20	2.32%	—	article, station, number, year, state, route, highway, time, road, points, date, railway, britannica, ship, include, information, eb, county, class, official

**Table F.2.** Top pages within each cluster, by eigenvector centrality.

Rank	Cluster Name	Top Pages
1	Article Quality	Neutral_point_of_view; Verifiability; Identifying_reliable_sources; What_Wikipedia_is_not; Biographies_of_living_persons; No_original_research; Citing_sources
2	Collaboration	Consensus; Policies_and_guidelines; Assume_good_faith; Dispute_resolution; Civility; Edit_warring; Talk_page_guidelines
3	Administrators	Administrators; Blocking_policy; Arbitration_Committee; Vandalism; User_pages; Sock_puppetry; User_access_levels
4	Formatting Articles	Redirect; Article_titles; Disambiguation; Manual_of_Style; Namespace; What_is_an_article?; Categorization
5	Content Policies	Copyrights; Copyright_violations; Non-free_content; Image_use_policy; General_disclaimer; Non-Wikipedia_disclaimers; Substitution
6	Wiki-larping	Citation_needed; Wikibreak; WikiGnome; Wikipediaholic; Talk_page_stalker; Wikipedia_is_a_volunteer_service; WikiDragon
7	Page Templates	Overlink_crisis; Pruning_article_revisions; Disinfoboxes; Thinking_outside_the_infobox; Advanced_template_coding; Advanced_article_editing; Advanced_footnote_formatting
8	Experts and Credentials	Expert_editors; Honesty; Expert_retention; Randy_in_Boise; Ten_Simple_Rules_for_Editing_Wikipedia; Conflicts_of_interest_(medicine); There_is_no_credential_policy
9	Humor	Silly_Things; Rules_for_Fools; April_Fools; April_Fool's_Main_Page; Unusual_articles; Yet_more_Best_of_BJAODN; Best_of_BJAODN

## References

1. Sherif, M. *The Psychology of Social Norms*; Harper: New York, NY, USA, 1936.
2. Durkheim, E. *The Rules of Sociological Method*; Free Press: New York, NY, USA, 1938.
3. Akerlof, G. The economics of caste and of the rat race and other woeful tales. *Quarterly Journal of Economics* **1976**, *90*, 599–617.
4. Geertz, C. Thick description: Toward an interpretive theory of culture. In *Readings in the Philosophy of Social Science*; Martin, M., McIntyre, L.C., Eds.; MIT Press: Cambridge, MA, USA, 1994; pp. 213–231.
5. Ellickson, R.C.; Ellickson, R.C. *Order without Law: How Neighbors Settle Disputes*; Harvard University Press: Cambridge, MA, USA, 2009.
6. Bowles, S. *Microeconomics: Behavior, Institutions, and Evolution*; Princeton University Press: Princeton, NJ, USA, 2009.
7. Simon, H.A. A formal theory of the employment relationship. *Econometrica* **1951**, *19*, 293–305.
8. Brennan, G.; Buchanan, J.M. *The reason of rules*; Cambridge University Press: Cambridge, United Kingdom, 2008.
9. Tyler, T.R. Psychological perspectives on legitimacy and legitimation. *Annual Review of Psychology* **2006**, *57*, 375–400.
10. Tyler, T.R.; Fagan, J. Legitimacy and cooperation: Why do people help the police fight crime in their communities. *Ohio State Journal of Criminal Law* **2008**, *6*, 231.
11. Elias, N. *The Civilizing Process: Sociogenetic and Psychogenetic Investigations*, 2nd ed.; Dunning, E., Goudsblom, J., Mennell, S., Eds.; Wiley: New York, NY, USA, 2000.
12. Pinker, S. *The Better Angels of Our Nature: Why Violence Has Declined*; Penguin Group: New York, NY, USA, 2011.
13. Klingenstein, S.; Hitchcock, T.; DeDeo, S. The civilizing process in London's Old Bailey. *Proceedings of the National Academy of Sciences* **2014**, *111*, 9419–9424.
14. Ehrlich, P.R.; Levin, S.A. The evolution of norms. *PLoS Biology* **2005**, *3*, 943.
15. Ostrom, E.; Hess, C. A framework for analyzing the knowledge commons. In *Understanding Knowledge as a Commons*; Hess, C., Ostrom, E., Eds.; MIT Press: Cambridge, MA, USA, 2006.
16. Benkler, Y. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*; Yale University Press: New Haven, CT, USA, 2006.
17. Bollier, D. The growth of the commons paradigm. In *Understanding Knowledge as a Commons*; Hess, C., Ostrom, E., Eds.; MIT Press: Cambridge, MA, USA, 2006.
18. Frischmann, B.; Madison, M.; Strandburg, K. *Governing Knowledge Commons*; Oxford University Press: Oxford, UK, 2014.
19. Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*; Cambridge University Press: Cambridge, UK, 1990.
20. Hess, C.; Ostrom, E. *Understanding Knowledge as a Commons: From Theory to Practice*; MIT Press: Cambridge, MA, USA, 2011.
21. West, J.; Lakhani, K.R. Getting clear about communities in open innovation. *Industry and Innovation* **2008**, *15*, 223–231.
22. O'Mahony, S. The governance of open source initiatives: What does it mean to be community managed? *Journal of Management & Governance* **2007**, *11*, 139–150.
23. Beschastnikh, I.; Kriplean, T.; McDonald, D.W. Wikipedian self-governance in action: Motivating the policy lens. In Proceedings of the ICWSM, Seattle, WA, 30 March–2 April 2008.
24. March, J.G.; Schulz, M.; Zhou, X. *The Dynamics of Rules: Change in Written Organizational Codes*; Stanford University Press: Palo Alto, CA, USA, 2000.
25. Butler, B.; Joyce, E.; Pike, J. Don't look now, but we've created a bureaucracy: The nature and roles of policies and rules in wikipedia. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Florence, Italy, 5–10 April 2008; ACM: New York, NY, USA, 2008; pp. 1101–1110.
26. Schneider, J.; Passant, A.; Breslin, J. A qualitative and quantitative analysis of how Wikipedia talk pages are used. In Proceedings of the 2010 ACM Conference on Web Science, Raleigh, NC, USA, 26–27 April 2010; ACM: New York, NY, USA, 2010.



27. Kriplean, T.; Beschastnikh, I.; McDonald, D.W.; Golder, S.A. Community, consensus, coercion, control: CS\*W or how policy mediates mass participation. In Proceedings of the 2007 International ACM Conference on Supporting Group Work, Sanibel Island, FL, USA, 4–7 November 2007; ACM: New York, NY, USA, 2007; pp. 167–176.
28. Park, H.W.; Thelwall, M. Hyperlink analyses of the World Wide Web: A review. *Journal of Computer-Mediated Communication* **2003**, *8*, 4.
29. Gonzalez-Bailon, S. Opening the black box of link formation: Social factors underlying the structure of the web. *Social Networks* **2009**, *31*, 271–280.
30. Strube, M.; Ponzetto, S.P. WikiRelate! Computing semantic relatedness using Wikipedia. In Proceedings of the AAAI 21st National Conference on Artificial Intelligence, Boston, MA, USA, 16–20 July 2006; Volume 6, pp. 1419–1424.
31. Witten, I.; Milne, D. An effective, low-cost measure of semantic relatedness obtained from Wikipedia links. In Proceeding of the AAAI Workshop on Wikipedia and Artificial Intelligence: An Evolving Synergy, 13 July, 2008 Chicago, Illinois, USA; AAAI Press: Menlo Park, CA, USA, 2008; pp. 25–30.
32. Bellomi, F.; Bonato, R. Network analysis for Wikipedia. In Proceedings of Wikimania, Frankfurt am Main, Germany, 4–8 August 2005.
33. Lizorkin, D.; Medelyan, O.; Grineva, M. Analysis of community structure in Wikipedia. In Proceedings of the 18th International Conference on World Wide Web, Madrid, Spain, 20–24 April 2009; ACM: New York, NY, USA, 2009; pp. 1221–1222.
34. Fowler, J.H.; Jeon, S. The authority of Supreme Court precedent. *Social Networks* **2008**, *30*, 16–30.
35. Walsh, D.J. On the meaning and pattern of legal citations: Evidence from state wrongful discharge precedent cases. *Law and Society Review* **1997**, *31*, 337–361.
36. Caldeira, G.A. The transmission of legal precedent: A study of state Supreme Courts. *American Political Science Review* **1985**, *79*, 178–194.
37. Henrich, J.; Boyd, R.; Richerson, P.J. Five misunderstandings about cultural evolution. *Human Nature* **2008**, *19*, 119–137.
38. Shirky, C. *Here Comes Everybody: The Power of Organizing without Organizations*; Penguin: New York, NY, USA, 2008.
39. Konieczny, P. Governance, Organization, and Democracy on the Internet: The Iron Law and the Evolution of Wikipedia. *Sociological Forum* **2009**, *24*, 162–192.
40. Konieczny, P. Adhocratic governance in the Internet age: A case of Wikipedia. *Journal of Information Technology & Politics* **2010**, *7*, 263–283.
41. Meyer, J.W.; Rowan, B. Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology* **1977**, *83*, 340–363.
42. Open Data for the paper the Evolution of Wikipedia's Norm Network. Available online: [https://bit.ly/wiki\\_norm](https://bit.ly/wiki_norm). (accessed on 21 August 2015).
43. Morgan, J.T.; Zachry, M. Negotiating with angry mastodons: The wikipedia policy environment as genre ecology. In Proceedings of the 16th ACM International Conference on Supporting Group Work, Sanibel, FL, USA, 7–10 November 2010; ACM: New York, NY, USA, 2010; pp. 165–168.
44. Template:Policy. Available online: <https://en.wikipedia.org/wiki/Template:Policy>. (accessed on 17 April 2016).
45. Template:Guideline. Available online: <https://en.wikipedia.org/wiki/Template:Guideline>. (accessed on 17 April 2016).
46. Template:Essay. Available online: <https://en.wikipedia.org/wiki/Template:Essay>. (accessed on 17 April 2016).
47. Template:Proposed. Available online: <https://en.wikipedia.org/wiki/Template:Proposed>. (accessed on 17 April 2016).
48. Cohen, J. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* **1960**, *20*, 37–46.
49. Halfaker, A.; Geiger, R.S.; Morgan, J.T.; Riedl, J. The rise and decline of an open collaboration system: How Wikipedia's reaction to popularity is causing its decline. *American Behavioral Scientist* **2013**, *57*, 664–688.
50. Wikipedia Statistics: Active Wikipedians. Available online: <https://stats.wikimedia.org/EN/TablesWikipediansEditsGt5.htm> (accessed on 21 August 2015).

51. Brush, E.R.; Krakauer, D.C.; Flack, J.C. A family of algorithms for computing consensus about node state from network data. *PLoS Computational Biology* **2013**, *9*, e1003109.
52. StatsGrok. Available online: <http://stats.grok.se>. (accessed on 17 April 2016). Data from service created by Domas Mituzas, visualized by Wikipedia User Henrik.
53. Blei, D.M.; Ng, A.Y.; Jordan, M.I. Latent dirichlet allocation. *Journal of Machine Learning Research* **2003**, *3*, 993–1022.
54. DeDeo, S.; Hawkins, R.X.; Klingenstein, S.; Hitchcock, T. Bootstrap methods for the empirical study of decision-making and information flows in social systems. *Entropy* **2013**, *15*, 2246–2276.
55. Blondel, V.D.; Guillaume, J.L.; Lambiotte, R.; Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **2008**, *2008*, P10008.
56. Yan, E.; Ding, Y. Scholarly network similarities: How bibliographic coupling networks, citation networks, cocitation networks, topical networks, coauthorship networks, and cword networks relate to each other. *Journal of the American Society for Information Science and Technology* **2012**, *63*, 1313–1326.
57. J. Richard Landis, G.G.K. The Measurement of Observer Agreement for Categorical Data. *Biometrics* **1977**, *33*, 159–174.
58. Jacomy, M.; Venturini, T.; Heymann, S.; Bastian, M. ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLoS ONE* **2014**, *9*, e98679.
59. Merton, R.K. The Matthew effect in science. *Science* **1968**, *159*, 56–63.
60. Shaw, A.; Hill, B.M. Laboratories of oligarchy? How the Iron Law extends to peer production. *Journal of Communication* **2014**, *64*, 215–238.
61. Forte, A.; Larco, V.; Bruckman, A. Decentralization in Wikipedia governance. *Journal of Management Information Systems* **2009**, *26*, 49–72.
62. Cohen, J. A power primer. *Psychological Bulletin* **1992**, *112*, 155.



© 2016 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons by Attribution (CC-BY) license (<http://creativecommons.org/licenses/by/4.0/>).